

Communication-Avoiding Parallel Algorithms for Solving Triangular Systems of Linear Equations

Tobias Wicky

Department of Computer Science
ETH Zurich
Zurich, Switzerland
twicki@ethz.ch

Edgar Solomonik

Department of Computer Science
University of Illinois at Urbana-Champaign
Urbana, IL, USA
solomon2@illinois.edu

Torsten Hoefer

Department of Computer Science
ETH Zurich
Zurich, Switzerland
htor@inf.ethz.ch

Abstract—We present a new parallel algorithm for solving triangular systems with multiple right hand sides (TRSM). TRSM is used extensively in numerical linear algebra computations, both to solve triangular linear systems of equations as well as to compute factorizations with triangular matrices, such as Cholesky, LU, and QR. Our algorithm achieves better theoretical scalability than known alternatives, while maintaining numerical stability, via selective use of triangular matrix inversion. We leverage the fact that triangular inversion and matrix multiplication are more parallelizable than the standard TRSM algorithm. By only inverting triangular blocks along the diagonal of the initial matrix, we generalize the usual way of TRSM computation and the full matrix inversion approach. This flexibility leads to an efficient algorithm for any ratio of the number of right hand sides to the triangular matrix dimension. We provide a detailed communication cost analysis for our algorithm as well as for the recursive triangular matrix inversion. This cost analysis makes it possible to determine optimal block sizes and processor grids a priori. Relative to the best known algorithms for TRSM, our approach can require asymptotically fewer messages, while performing optimal amounts of communication and computation.

Keywords—TRSM, communication cost, 3D algorithms

I. INTRODUCTION

Triangular solve for multiple right hand sides (TRSM) is a crucial subroutine in many numerical linear algebra algorithms, such as LU and Cholesky factorizations [1], [2]. Moreover, it is used to solve linear systems of equations once the equation matrix is decomposed using any factorization involving a triangular matrix factor. We consider TRSM for dense linear equations in the matrix form,

$$L \cdot X = B,$$

where $L \in \mathbb{R}^{n \times n}$ is a lower-triangular matrix while $B \in \mathbb{R}^{n \times k}$, $X \in \mathbb{R}^{n \times k}$ are dense matrices. We study the communication cost complexity of two variants of the TRSM algorithm for parallel execution on p processors with unbounded memory. We first present an adaptation of a known recursive scheme for TRSM [3] along with a complete communication cost analysis. Then, we demonstrate a new algorithm that uses triangular inversion of matrix blocks to reduce the synchronization cost over known schemes,

while preserving optimal communication and computation costs. Careful choice of algorithmic parameters, allows us to achieve better asymptotic complexity for a large (and most important) range of input configurations.

Our TRSM algorithm leverages matrix multiplication and triangular matrix inversion as primitives. We provide a communication-efficient parallel matrix multiplication algorithm that starts from a 2D cyclic distribution, a modest enhancement to existing approaches. Triangular matrix inversion provides the key ingredient to the lower synchronization cost in our TRSM algorithm. Unlike general matrix inversion, triangular inversion is numerically stable [4] and can be done with relatively few synchronizations. We present a known parallel approach for triangular inversion and provide the first communication cost analysis thereof.

We invert triangular diagonal blocks of the L matrix at the start of our TRSM algorithm, increasing the computational granularity of the main part of the solver. Inverting blocks, rather than the whole matrix, also allows our algorithm to be work-efficient in cases when the number of right-hand sides is smaller than the matrix dimension. We formulate the algorithm in an iterative, rather than a recursive manner, avoiding overheads incurred by the known parallel recursive TRSM approach. This innovation reduces communication cost by a factor of $\Theta(\log(p))$ relative to the recursive algorithm when the number of right-hand sides is relatively small. At the same time, across a large range of input parameters, we achieve a synchronization cost improvement over the recursive approach, by a factor of $\Theta\left(\left(\frac{n}{k}\right)^{1/6} p^{2/3}\right)$.

II. PRELIMINARIES

A. Execution Time Model

The model we use to calculate the parallel execution time of an algorithm along its critical path is the $\alpha - \beta - \gamma$ model. It describes the total execution time of the algorithm T in terms of the floating point operations (flop) count F , the bandwidth W (number of words of data sent and received)

and the latency S (number of messages sent and received) along the critical path [5] in the following fashion:

$$T = \alpha \cdot S + \beta \cdot W + \gamma \cdot F.$$

For all the terms we only show the leading order cost in terms of n, k and p . We assume that every processor can send and receive one message at a time in point to point communication. We do not place constraints on the local memory size.

B. Notation

For brevity, we will sometimes omit specification of the logarithm base, using \log to denote \log_2 . We will make frequent use of the unit step function

$$\mathbb{1}_x = \begin{cases} 1 & : x > 1 \\ 0 & : x \leq 1. \end{cases}$$

We will denote processors in a 3D processor grid as $\Pi(x, y, z)$ where $x, y \in [1, \dots, p_1]$ and $z \in [1, \dots, p_2]$.

To refer to successive elements or blocks, we will use the colon notation where $i : j = [i, i+1, \dots, j-1]$ $i < j$. To refer to strided sets of elements or blocks, we will write

$$i : k : j = [i, i+k, i+2k, \dots, i+\alpha k] \quad \max \alpha \text{ s.t. } \alpha k < j.$$

The colon notation can be applied to a list and should there be considered element-wise. To apply operations on a set of processors, we use the \circ notation such that $\Pi(x, \circ, z) = \Pi(x, 1 : p_1, z)$ denotes a (1-dimensional) grid of processors in the y -dimension.

C. Previous Work

We now cover necessary existing building blocks (e.g. collective communication routines) and previous work. In particular, we overview related results on communication cost of matrix multiplication and triangular solves.

1) Collective Communication: The most important collective communication methods that we use to do three dimensional matrix multiplications are broadcasts, (all)gathers, (all)-reductions as well as reduce-scatter. In [6], Chan et al. present a way to perform reduction, allreduction and broadcast via allgather, scatter, gather and reduce-scatter. The latter set of collectives can be done using recursive doubling (butterfly algorithms) [6], [7] for a power of two number of processors. If we have a non-power of two number of processors, the algorithm described in [8] can be used. For simplicity, we do not consider reduction and broadcast algorithms that can achieve a factor of two less in cost in specific regimes of α and β [9]. If we use butterfly methods, the cost of an allgather of n words among p processors is

$$T_{\text{allgather}}(n, p) = \alpha \cdot \log p + \beta \cdot n \mathbb{1}_p.$$

Scatter and gather also have the same cost [7],

$$\begin{aligned} T_{\text{scatter}}(n, p) &= \alpha \cdot \log p + \beta \cdot n \mathbb{1}_p, \\ T_{\text{gather}}(n, p) &= \alpha \cdot \log p + \beta \cdot n \mathbb{1}_p. \end{aligned}$$

Reduce-scatter uses the exact same communication-path and therefore has same cost except for the additional overhead of local computation,

$$T_{\text{reduce-scatter}}(n, p) = \alpha \cdot \log p + \beta \cdot n \mathbb{1}_p + \gamma \cdot n \mathbb{1}_p.$$

The combination of the algorithms gives us the following costs for reduction, allreduction, and broadcast:

$$\begin{aligned} T_{\text{reduction}}(n, p) &= \alpha \cdot 2 \log p + \beta \cdot 2n \mathbb{1}_p + \gamma \cdot n \mathbb{1}_p, \\ T_{\text{allreduction}}(n, p) &= \alpha \cdot 2 \log p + \beta \cdot 2n \mathbb{1}_p + \gamma \cdot n \mathbb{1}_p, \\ T_{\text{broadcast}}(n, p) &= \alpha \cdot 2 \log p + \beta \cdot 2n \mathbb{1}_p. \end{aligned}$$

The cost of an all-to-all among p processors is

$$T_{\text{alltoall}}(n, p) = \alpha \cdot \log(p) + \beta \cdot \frac{n \log p}{2}.$$

2) Matrix Multiplication: Communication-efficient parallel algorithms for matrix multiplication have been analyzed extensively [10]–[15]. In [16], Demmel et al. present algorithms that are asymptotically optimal for matrix multiplication of arbitrary (potentially non square) matrices. If we neglect the memory terms, their work shows that matrix multiplication can be done with the following costs.

Bandwidth: When multiplying a matrix A that is of dimension $n \times n$ with a matrix B of dimensions $n \times k$ with p processors, we obtain an asymptotic bandwidth cost of

$$W_{\text{MM}}(n, k, p) = \begin{cases} \mathcal{O}\left(\frac{nk}{\sqrt{p}}\right) & n > k \cdot \sqrt{p} \\ \mathcal{O}\left(\left(\frac{n^2 k}{p}\right)^{2/3}\right) & k/p \leq n \leq k \cdot \sqrt{p} \\ \mathcal{O}(n^2) & n < k/p. \end{cases}$$

We refer to the first of the cases of W_{MM} , as the case of two large dimensions, here the matrix A is much larger than the right hand side B , when the best way of performing a matrix multiplication is to use a two dimensional layout for the processor grid. The second case, three large dimensions, has matrices A and B of approximately the same size. A three dimensional grid layout is optimal here. And the third case, one large dimension, is the case where the right hand side B is larger than the triangular matrix A , the best way to do a matrix multiplication is to use a one dimensional layout for the processor grid.

Latency: Assuming unlimited memory, matrix multiplication as presented in [16] can be done with a latency of

$$S_{\text{MM}}(p) = \mathcal{O}(\log(p)).$$

Flop Cost: Matrix multiplication takes $\mathcal{O}(n^2 k)$ flops, which can be divided on p processors and therefore we have

$$F_{\text{MM}}(n, k, p) = \mathcal{O}\left(\frac{n^2 k}{p}\right).$$

Previous Analysis: For the case where $k = n$ the bandwidth analysis of a general matrix multiplication goes back to what is presented in [17]. Aggarwal et al. present a cost analysis in the LPRAM model. In that work, the

authors show that the same cost can also be achieved for the transitive closure problem that can be extended to the problem of doing an LU decomposition. The fact that these bandwidth costs can be obtained for the LU decomposition was later demonstrated by Tiskin [18]. He used the bulk synchronous parallel (BSP) execution time model. Since the dependencies in LU are more complicated than they are for TRSM, we also expect TRSM to be able to have the same asymptotic bandwidth and flop costs as a general matrix multiplication.

3) *Triangular Matrix Solve for Single Right Hand Sides:* Algorithms for the problem of triangular solve for a single right hand side (when X and B are vectors) have been well-studied. A communication-efficient parallel algorithm was given by Heath and Romine [19]. This parallel algorithm was later shown to be an optimal schedule in latency and bandwidth costs via lower bounds [5]. However, when X and B are matrices ($k > 1$), it is possible to achieve significantly lower communication costs relative to the amount of computation required.

4) *Recursive Triangular Matrix Solve for Multiple Right Hand Sides:* A recursive approach of solving the TRSM-problem was presented in the work of Elmroth et al. [3]. We can subdivide the initial problem, $L \cdot X = B$ as $L \cdot [X_1 \ X_2] = [B_1 \ B_2]$ which yields two independent subproblems:

$$L \cdot X_1 = B_1, \quad L \cdot X_2 = B_2,$$

which can be solved in parallel.

The other, dependent splitting proposed divides the triangular matrix, again yielding two subtasks that have to be solved one at a time

$$\begin{bmatrix} L_{11} & \\ L_{12} & L_{22} \end{bmatrix} \cdot \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} = \begin{bmatrix} B_1 \\ B_2 \end{bmatrix},$$

where we obtain the dependent subproblems:

$$L_{11} \cdot X_1 = B_1 \quad \text{and} \quad L_{22} \cdot X_2 = B_2 - L_{12} \cdot X_1.$$

These problems are dependent as we need the solution X_1 to solve the second problem. Parallel algorithms with 3D processor grids can reduce the communication cost in an analogous fashion to matrix multiplication. Irony and Toledo [20] presented the first parallelization of the recursive TRSM algorithm with a 3D processor grid. They demonstrated that the communication volume of their parallelization is $\mathcal{O}(nkp^{1/3} + n^2p^{1/3})$. Thus each processor communicates $\mathcal{O}((nk + n^2)/p^{2/3})$ elements, which is asymptotically equal to $W_{\text{MM}}(n, k, p)$ when $k = \Theta(n)$. However, they did not provide a bound on the latency cost nor on the communication bandwidth cost along the critical path, so it is unclear to what extent the communication volume is load-balanced. Lipshitz [21] provides an analysis of the recursive TRSM algorithm in the same communication cost model as used in this paper. For the case of $k = n$, his analysis demonstrates

$$T_{\text{TRSM-L}}(n, n, p) = \mathcal{O}(n^3/p \cdot \gamma + n^2/p^{2/3} \cdot \beta + p^{2/3} \cdot \alpha).$$

For some choices of algorithmic parameters, the analysis in [21] should lead to the bandwidth cost,

$$W_{\text{TRSM-L}}(n, k, p) = \mathcal{O}\left(\left(\frac{n^2k}{p}\right)^{2/3}\right),$$

which is as good as matrix multiplication, $W_{\text{MM}}(n, k, p)$. However, it is unclear how to choose the parameters of the formulation in [21] to minimize latency cost for general n, k . Prior to presenting our main contribution (an inversion-based TRSM algorithm with a lower latency cost), we provide a simpler form of the recursive TRSM algorithm and its asymptotic cost complexity. We start by presenting a subroutine for 3D matrix multiplication which operates from a starting 2D distribution, simplifying the subsequent presentation of TRSM algorithms.

III. MATRIX MULTIPLICATION

We present an algorithm for 3D matrix multiplication [10]–[15] that works efficiently given input matrices distributed cyclically on a 2D processor grid. The algorithm is well-suited for the purposes of analyzing TRSM algorithms. We define the algorithm using a $p_1 \times \sqrt{p_2} \times p_1 \times \sqrt{p_2}$ processor grid, where $\sqrt{p} = p_1\sqrt{p_2}$ in order to provide well-defined transitions from a distribution on a $\sqrt{p} \times \sqrt{p}$ processor grid to faces of a 3D $p_1 \times p_1 \times p_2$ processor grid. The latter 3D processor grid is being used implicitly in our construction. The algorithm assumes divisibility among p, p_1, p_2 and $\sqrt{p_2}$.

$$B = \mathbf{MM}(L, X, \Pi_{2D}, n, k, p, p_1, p_2)$$

Require:

The processor grid Π_{2D} has dimensions $\sqrt{p} \times \sqrt{p}$
 L is an $n \times n$ matrix, distributed on Π_{2D} in a cyclic layout, so $\Pi_{2D}(x, y)$ owns $L[x, y] \in \mathbb{R}^{\frac{n}{\sqrt{p}} \times \frac{n}{\sqrt{p}}}$ such that
 $L[x, y](i, j) = L(i\sqrt{p} + x, j\sqrt{p} + y)$.
 X is a dense $n \times k$ matrix is distributed cyclically so that
 $\Pi_{2D}(x, y)$ owns $X[x, y] \in \mathbb{R}^{\frac{n}{\sqrt{p}} \times \frac{k}{\sqrt{p}}}$

- 1: Define a $p_1 \times \sqrt{p_2} \times p_1 \times \sqrt{p_2}$ processor grid Π_{4D} , such that $\Pi_{4D}(x_1, x_2, y_1, y_2) = \Pi_{2D}(x_1 + p_1x_2, y_1 + p_2y_2)$ owns blocks $L[x_1, x_2, y_1, y_2]$ and $X[x_1, x_2, y_1, y_2]$
- 2: $L'[x_1, y_1] = \mathbf{Allgather}(L[x_1, \circ, y_1, \circ], \Pi_{4D}(x_1, \circ, y_1, \circ))$
- 3: $X'[x_1, y_1, x_2, y_2] = \mathbf{Transpose}(X[x_1, x_2, y_1, y_2], \Pi_{4D}(x_1, x_2, y_1, y_2), x_2, y_1)$
- 4: $X''[y_1, x_1, x_2, y_2] = \mathbf{Transpose}(X'[x_1, y_1, x_2, y_2], \Pi_{4D}(x_1, x_2, y_1, y_2), x_1, y_1)$
- 5: $X'''[y_1, x_2, y_2] = \mathbf{Allgather}(X''[y_1, \circ, x_2, y_2], \Pi_{4D}(\circ, x_2, y_1, y_2))$
- 6: $\Pi_{4D}(x_1, x_2, y_1, y_2) : B''[x_1, y_1, x_2, y_2] = L'[x_1, y_1] \cdot X'''[y_1, x_2, y_2]$
- 7: $B'[x_1, y_1, x_2, y_2] = \mathbf{Scatter-reduce}(B''[x_1, \circ, x_2, y_2], \Pi_{4D}(x_1, x_2, \circ, y_2))$
- 8: $B[x_1, x_2, y_1, y_2] = \mathbf{Transpose}(B'[x_1, y_1, x_2, y_2], \Pi_{4D}(x_1, x_2, y_1, y_2), x_2, y_1)$

Ensure:

$B = LX$ is distributed the same way as X

A. Cost Analysis of the 3D Matrix Multiplication Algorithm

We analyze the algorithm with account for constant factors in the key leading order costs. The communication costs incurred at line 2, 5, and 7 correspond to the cost of respective collectives, given in Section II-C1.

The transpose on line 4 always occurs on a square processor grid, and so involves only one send and receive of a block. The transposes on line 3 and line 8 are transposes on 2D grids of $p_1 \times \sqrt{p_2}$ processors, with each processor owning nk/p elements. The cost of these transposes is no greater than an all-to-all among \sqrt{p} processors, which can be done with cost $\mathcal{O}(\alpha \cdot \log(p) + \beta \cdot nk \log(p)/p)$. We consider only the asymptotic cost of this transpose, since it will be low order so long as $p_1 \gg 1$. Based on the above arguments, the cost for MM is given line-by-line in the following table.

Line 2	$\alpha \cdot \log(p_2) + \beta \cdot \frac{n^2}{p_1^2} \mathbb{1}_{p_2}$
Line 3	$\mathcal{O}\left(\alpha \cdot \log(p) + \beta \cdot \frac{nk \log(p)}{p}\right)$
Line 5	$\alpha \cdot \log(p_1) + \beta \cdot \frac{nk}{p_1 p_2}$
Line 4	$\alpha + \beta \cdot \frac{nk}{p}$
Line 6	$\gamma \cdot \frac{n^2 k}{p}$
Line 7	$\alpha \cdot \log(p_1) + (\beta + \gamma) \cdot \frac{nk}{p_1 p_2}$
Line 8	$\alpha \cdot l + \beta \cdot \frac{nk l}{p}$

To leading order, the cost of MM is given by

$$T_{\text{MM}}(n, k, p, p_1, p_2) = \beta \cdot \left(\frac{n^2}{p_1^2} \mathbb{1}_{p_2} + \frac{2nk}{p_1 p_2} \right) + \gamma \cdot \frac{n^2 k}{p} + \mathcal{O}\left(\alpha \cdot \log(p) + \beta \cdot \frac{nk \log(p)}{p}\right).$$

The last communication cost term (due to the rectangular grid transpose) is only of leading order when $p_1 \approx \log(p)$. A square processor grid is not a good initial/final layout for X and B in this case, and the problem would be addressed by choosing an alternative one. We disregard this issue, because we will use the algorithm only for $n \geq k$.

IV. RECURSIVE TRSM

We now provide a recursive TRSM algorithm for solving $LX = B$ using the techniques covered in Section II-C4. Our algorithm works recursively on a $p_r \times p_c$ processor grid. We will define the processor grid to be square ($p_r = p_c$) when $n \geq k$, but rectangular ($p_r < p_c$) when $n < k$. So long as $p < k/n$, we will choose $p_c = (k/n)p_r$. This strategy implies the largest of the matrices L and B will be partitioned initially so each cyclically-selected block is close to square. The algorithm starts by partitioning the processor grid into p_c/p_r square grids, if $p_c > p_r$, replicating the matrix L and computing a subset of kp_r/p_c columns of X on each. Then the algorithm partitions L into $n/2 \times n/2$ blocks recursively, executing subproblems with all processors. At a given threshold, n_0 , the algorithm stops recursing, gathers

L onto all processors, and computes a subset of columns of X with each processor.

$$X = \text{Rec-TRSM}(L, B, \Pi_{2D}, n, k, p_r, p_c, n_0)$$

Require:

L is a lower triangular $n \times n$ matrix, distributed on $p_r \times p_c$ in a cyclic layout, so $\Pi_{2D}(x, y)$ owns $L[x, y] \in \mathbb{R}^{\frac{n}{p_r} \times \frac{n}{p_c}}$ such that $L[x, y](i, j) = L(ip_r + x, jp_c + y)$.

B is a dense $n \times k$ matrix is distributed cyclically so that $\Pi_{2D}(x, y)$ owns $X[x, y] \in \mathbb{R}^{\frac{n}{p_r} \times \frac{k}{p_c}}$

- 1: **if** $p_r = qp_c$ and $q > 1$ **then**
 - 2: Define $p_r \times p_r \times q$ processor grid Π_{3D} , such that $\Pi_{3D}(x, y, z) = \Pi_{2D}(x_1, y + p_r z)$ owns blocks $L[x, y, z]$, and $B[x, y, z]$
 - 3: $L'[x, y] = \text{Allgather}(L[x, y, \circ], \Pi_{3D}(x, y, \circ))$
 - 4: $X[\circ, \circ, z] = \text{Rec-TRSM}(L'[\circ, \circ], B[\circ, \circ, z], \Pi_{2D}(\circ, \circ, z), n, k/q, p_r, p_r, n_0)$
 - 5: **else if** $n \leq n_0$ or $p_r = p_c = 1$ **then**
 - 6: $L = \text{Allgather}(L[\circ, \circ], \Pi_{2D}(\circ, \circ))$
 - 7: $B[x + yp_r] = \text{AllToAll}(B[\circ, y], \Pi_{2D}(\circ, y))$
 - 8: $\Pi_{2D}(x, y) : X[x + yp_r] = L^{-1} B[x + yp_r]$
 - 9: $X[x, y] = \text{AllToAll}(B[x + \circ p_r], \Pi_{2D}(\circ, y))$
 - 10: **else**
 - 11: Partition $L = \begin{bmatrix} L_{11} & 0 \\ L_{21} & L_{22} \end{bmatrix}$ so $L_{ij} \in \mathbb{R}^{\frac{n}{2} \times \frac{n}{2}}$
 - 12: Partition $B = \begin{bmatrix} B_1 \\ B_2 \end{bmatrix}$, $X = \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}$, so $B_i, X_i \in \mathbb{R}^{\frac{n}{2} \times k}$
 - 13: $X_1 = \text{Rec-TRSM}(L_{11}, B_1, \Pi_{2D}, n/2, k, p, p_r, p_c, n_0)$.
 - 14: $B'_2 = B_2 - \text{MM}(L_{21}, X_1, \Pi_{2D}, n/2, k, p, p_r, p_c, n_0)$
 $\quad \quad \quad p_r^{1/3}(n/k)^{1/3}, p_r^{1/3}(n/k)^{2/3}$.
 - 15: $X_2 = \text{Rec-TRSM}(L_{22}, B'_2, \Pi_{2D}, n/2, k, p, p_r, p_c, n_0)$.
 - 16: **end if**
- Ensure:**
 $X = L^{-1}B$ is distributed on Π_{2D} in the same way as B
-

A. Cost Analysis of the Recursive Algorithm

We select $p_c = \max(\sqrt{p}, \min(p, \sqrt{pk/n}))$ and $p_r = p/p_c = \min(\sqrt{p}, \max(1, \sqrt{pn/k}))$. The cost of the allgather on line 3 is

$$T_{\text{part-cols}}(n, p_r) = \mathcal{O}\left(\beta \cdot \frac{n^2}{p_r^2} + \alpha \cdot \log(p)\right),$$

since each $L'[x, y] \in \mathbb{R}^{n/p_r \times n/p_r}$ and is lower triangular. Once we have a square processor grid, we partition L , yielding the recurrence,

$$T_{\text{RT}}(n, k, p, n_0) = T_{\text{MM}}\left(n/2, k, p, p^{1/3}\left(\frac{n}{k}\right)^{1/3}, p^{1/3}\left(\frac{n}{k}\right)^{2/3}\right) + 2T_{\text{RT}}(n/2, k, p, n_0).$$

We now derive the cost of the algorithm for different relations between n , k , and p , as in the expression for T_{MM} .

One large dimension: When $n < k/p$, we have $p_r = 1$ and $p_c = p$ and the first allgather will be the only communication, therefore,

$$T_{\text{RTID}}(n, k, p) = \mathcal{O}\left(\alpha \cdot \log(p) + \beta \cdot n^2 + \gamma \cdot \frac{n^2 k}{p}\right).$$

Two large dimensions: When $n > k\sqrt{p}$, we will select $p_r = p_c = \sqrt{p}$ and the column partitioning of B is not performed. In this case, the MM algorithm will always have $p_2 = 1$ (it will be 2D). For sufficiently large k , this leads us to the recurrence,

$$T_{\text{RT2D}}(n, k, p, n_0) = T_{\text{MM2D}}(n/2, k, p) + 2T_{\text{RT2D}}(n/2, k, p, n_0),$$

where

$$T_{\text{MM2D}}(n, k, p) = \mathcal{O}\left(\alpha \cdot \log(p) + \beta \cdot \frac{nk}{\sqrt{p}} + \gamma \cdot \frac{n^2k}{p}\right).$$

The bandwidth cost stays the same at every recursive level, while the computation cost decreases by a factor of 2. At the base case, we incur the cost $T_{\text{RTBC}}(n_0, k, p) =$

$$\mathcal{O}\left(\alpha \cdot \log(p) + \beta \cdot \left(n_0^2 + \frac{n_0k \log(p)}{p}\right) + \gamma \cdot \frac{n_0^2k}{p}\right)$$

We select $n_0 = \max(\sqrt{p}, n \log(p)/\sqrt{p})$, so $n/n_0 \leq \sqrt{p}/\log(p)$, which results in the overall cost,

$$\begin{aligned} T_{\text{RT2D}}(n, k, p) &= \mathcal{O}\left(\alpha \cdot \log^2(p) + \beta \cdot \frac{nk \log(p)}{\sqrt{p}} + \gamma \cdot \frac{n^2k}{p}\right) \\ &+ \frac{\sqrt{p}}{\log(p)} T_{\text{RTBC}}(\max(\sqrt{p}, n/\sqrt{p}), k, p) \\ &= \mathcal{O}\left(\alpha \cdot \sqrt{p} + \beta \cdot \frac{nk \log(p)}{\sqrt{p}} + \gamma \cdot \frac{n^2k}{p}\right). \end{aligned}$$

The bandwidth cost above is suboptimal by a factor of $\mathcal{O}(\log(p))$. The overhead is due to the recursive algorithm re-broadcasting some of the same elements of L at every recursive level. We use an iterative approach for our subsequent TRSM algorithm to avoid this redundant communication.

Three large dimensions:

When $k/p < n < k/\sqrt{p}$, the algorithm partitions the columns of B initially then recursively partitions L . In particular, we select $\max(\sqrt{p}, \sqrt{pk/n})$ and $p_r = p/p_c = \min(\sqrt{p}, \sqrt{pn/k})$, so the first step partitions a rectangular processor grid into $\max(1, k/n)$ fewer processor grids. After the first step, which partitions B , we have independent subproblems with p_r processors. We now start recursively partitioning L , yielding the cost recurrence,

$$\begin{aligned} T_{\text{RT3D}}(n, k, p_r^2, n_0) &= T_{\text{MM3D}}(n/2, k, p_r^2) \\ &+ T_{\text{part-cols}}(n, p) \mathbb{1}_{k/n} + 2T_{\text{RT3D}}(n/2, k, p_r^2, n_0), \end{aligned}$$

where we are always employing the MM algorithm in the 3D regime by selecting $p_1 = p_r^{2/3}(n/l)^{1/3}$ and $p_2 = p_r^{2/3}(n/l)^{2/3}$ where $l = kp_r/p_c$. In this case, T_{MM} reduces to $T_{\text{MM3D}}(n, k, p) =$

$$\mathcal{O}\left(\alpha \cdot \log(p) + \beta \cdot \left(\left(\frac{n^2k}{p}\right)^{2/3} + \frac{nk \log(p)}{p}\right) + \gamma \cdot \frac{n^2k}{p}\right)$$

This gives to the cost recurrence, $T_{\text{RT3D}}(n, k, p_r^2, n_0) =$

$$\begin{aligned} &\mathcal{O}\left(\alpha \cdot \log(p) + \beta \cdot \left(\left(\frac{n^2k}{p_r^2}\right)^{2/3} + \frac{n^2}{p_r^2} \mathbb{1}_{\frac{k}{n}}\right.\right. \\ &\left.\left. + \frac{nk \log(p)}{p_r^2}\right) + \gamma \cdot \frac{n^2k}{p_r^2}\right) + 2T_{\text{RT3D}}\left(\frac{n}{2}, k, p_r^2, n_0\right), \end{aligned}$$



Figure 1. Inversion of a Triangular Matrix

where we can see that $\frac{nk \log(p)}{p_r^2} = \mathcal{O}((n^2k/p_r^2)^{2/3})$, since the initial partitioning will give $n \geq k$. Its also easy to see that the $\frac{n^2}{p_r^2} \mathbb{1}_{\frac{k}{n}} = \mathcal{O}((n^2k/p_r^2)^{2/3})$. With these simplifications,

$$\begin{aligned} T_{\text{RT3D}}(n, k, p_r^2, n_0) &= \mathcal{O}\left(\alpha \cdot \log(p) + \beta \cdot \left(\frac{n^2k}{p_r^2}\right)^{2/3}\right. \\ &\left. + \gamma \cdot \frac{n^2k}{p_r^2}\right) + 2T_{\text{RT3D}}\left(\frac{n}{2}, k, p_r^2, n_0\right). \end{aligned}$$

We observe that the bandwidth cost $\mathcal{O}((n^2k/p_r^2)^{2/3})$ decreases by a factor of $2^{1/3}$ at every recursive level, and the computation cost by a factor of 2. The base-case cost will be $T_{\text{RTBC}}(n_0, k, p_r^2)$. We select $n_0 = n^{1/3} \left(\frac{k}{p_r^2}\right)^{2/3}$, giving a total cost over all base cases of $\frac{n}{n_0} T_{\text{RTBC}}(n_0, k, p_r^2) =$

$$\begin{aligned} &\mathcal{O}\left(\alpha \cdot \frac{n}{n_0} \log(p) + \beta \cdot \left(nn_0 + \frac{nk \log(p)}{p_r^2}\right) + \gamma \cdot \frac{nn_0k}{p_r^2}\right) \\ &= \mathcal{O}\left(\alpha \cdot \left(\frac{np_r^2}{k}\right)^{2/3} \log(p) + \beta \cdot \left(\frac{n^2k}{p_r^2}\right)^{2/3} + \gamma \cdot \frac{n^{4/3}k^{5/3}}{p_r^{10/3}}\right). \end{aligned}$$

Therefore, the overall cost incurred on each square processor grid is $T_{\text{RT3D}}(n, k, p_r^2) =$

$$\mathcal{O}\left(\alpha \cdot \left(\frac{np_r^2}{k}\right)^{2/3} \log(p) + \beta \cdot \left(\frac{n^2k}{p_r^2}\right)^{2/3} + \gamma \cdot \frac{n^2k}{p_r^2}\right).$$

When $k \leq n$, we do not have a partitioning step and $p_r^2 = p$. Otherwise, we have $p_r^2 = np/k$ obtain the cost $T_{\text{RT3D}}(n, n, np/k) =$

$$\mathcal{O}\left(\alpha \cdot \left(\frac{np}{k}\right)^{2/3} \log(p) + \beta \cdot \left(\frac{n^2k}{p}\right)^{2/3} + \gamma \cdot \frac{n^2k}{p}\right),$$

which is the same as for the case $k \leq n$. For $n = k$, the 3D costs obtained above are the same as the most efficient algorithms for $n \times n$ LU factorization. In the subsequent sections, we show that a lower synchronization cost is achievable via selective use of triangular matrix inversion.

V. TRIANGULAR INVERSION

In this section, we derive the cost of inverting a lower triangular matrix $L \in \mathbb{R}^{n \times n}$ with a total of p processors. Since the matrix to be inverted is square, the dimensions of the processor grid Π should be the same in two dimensions leaving us with $\dim(\Pi) = p_1 \times p_1 \times p_2$, where $p = p_1^2 p_2$. We assume the initial matrix to be cyclically distributed on the subgrid $\Pi(\circ, \circ, 1)$.

A. Algorithmic Approach

In [22], a recursive method for inverting triangular matrices was presented. A recursive method for full inversion was presented in [23]. When applied to a triangular matrix, this method coincides with the method from [22]. The method uses the triangular structure of the initial matrix to calculate the inverse by subdividing the problem into two recursive matrix inversion calls, which are executed concurrently and then uses two matrix multiplications to complete the inversion (Fig. 1). Since the subproblems are independent, we want to split the processor grid such that two distinct sets of processors work on either subproblem. We chose the base-case condition to be that the grid is one-dimensional in the dimension of p_1 and we do redundant base-case calculations with this subgrid. For this section, we consider $p_2 > p_1$, a constraint that we will always fulfill whenever the method is called.

$$L^{-1} = \text{RecTriInv}(L, \Pi, p, p_1, p_2)$$

Require:

The processor grid Π has dimensions $\sqrt{p} \times \sqrt{p}$
 L is a lower triangular $n \times n$ matrix, distributed on Π in a cyclic layout, so $\Pi(x, y)$ owns $L[x, y] \in \mathbb{R}^{\frac{n}{\sqrt{p}} \times \frac{n}{\sqrt{p}}}$ such that $L[x, y](i, j) = L(i\sqrt{p} + x, j\sqrt{p} + y)$.

```

1: if  $p_1 = 1$  then
2:   AllToAll ( $L[\circ, y], \Pi(x, \circ)$ )
3:    $L^{-1} = \text{sequential inversion}(L)$ 
4: else
5:   Subdivide  $L$  into  $n/2 \times n/2$  blocks,
6:    $L = \begin{bmatrix} L_{11} & 0 \\ L_{21} & L_{22} \end{bmatrix}$ 
7:   Subdivide the processor grid  $\Pi = [\Pi_1, \Pi_2]$  such that
       $\dim(\Pi_1) = \dim(\Pi_2) = (\sqrt{p}/2 \times \sqrt{p}/2)$ 
8:   Redistribute ( $L_{11}, \Pi \rightarrow \Pi_1$ )
9:   Redistribute ( $L_{22}, \Pi \rightarrow \Pi_2$ )
10:   $L_{11}^{-1} = \text{Rec-Tri-Inv}(L_{11}, \Pi_1, p, p_1/2^{2/3}, p_2/2^{2/3})$ 
11:   $L_{22}^{-1} = \text{Rec-Tri-Inv}(L_{22}, \Pi_2, p, p_1/2^{2/3}, p_2/2^{2/3})$ 
12:   $L_{21}^{-1} = -\text{MM}(L_{22}^{-1}, L_{21}, \Pi, n, n, p, p_1, p_2)$ 
13:   $L_{21}^{-1} = \text{MM}(L_{21}^{-1}, L_{11}^{-1}, \Pi, n, n, p, p_1, p_2)$ 
14:  Assemble  $L^{-1}$  from the  $n/2 \times n/2$  blocks,
15:   $L^{-1} = \begin{bmatrix} L_{11}^{-1} & 0 \\ L_{21}^{-1} & L_{22}^{-1} \end{bmatrix}$ 
16: end if
```

Ensure:

$LL^{-1} = \mathbf{1}$ where L^{-1} is distributed the same way as L

B. Total Cost of Triangular Inversion

This recursive approach of inverting the matrix has a total cost of

$$T_{\text{RecTriInv}}(n, p_1, p_2) = 2T_{\text{MM}}(n/2, n/2, p_1, p_2) + T_{\text{RecTriInv}}(n/2, p_1/2^{1/3}, p_2/2^{1/3}) + T_{\text{redistr}}(n/2, p_1, p_2),$$

with a base case cost of

$$T_{\text{RecTriInv}}(n_0, 1, p_2) = \alpha \cdot 2 \log \left(\frac{p_2}{p_1} \right) + \beta \cdot 2n_0^2 + \gamma \cdot n_0^3.$$

The base case size will be $n_0 = \frac{n}{p_1^{3/2}}$ and therefore neither of the terms is of leading order. We observe that the bandwidth cost of the matrix multiplication $\mathcal{O}((n^3/p_1^2)^{2/3})$ decreases by a factor of $2^{4/9}$ at every recursive level, and the computation cost by a factor of 2. The redistribution process requires moving the matrices from a cyclic processor grid to a smaller cyclic processor grid, with the block each processor owns having a factor of $2^{1/3}$ more rows and columns. This redistribution is effectively an all-to-all between a larger and a smaller set of processors. We can get a concrete bound on the cost, by first performing an all-to-all to transition to a blocked layout (each processor owns contiguous blocks of the matrix). Then we can transition to a blocked layout on the smaller processor grid by scattering each block to at most 4 processors. Finally, we can perform an all-to-all on the smaller processor grid to transition from the blocked layout back to a cyclic one. The overall cost of these steps is $O(\alpha \cdot \log(p) + \beta \cdot n_0^2 \log(p)/p)$. This redistribution bandwidth cost is dominated by the cost of the matrix multiplication.

The total cost for the recursive inversion is

$$T_{\text{RecTriInv}}(n, p_1, p_2) = \beta \cdot \frac{2^{1/3}}{2^{1/3} - 1} \left(\frac{n^2}{8p_1^2} + \frac{n^2}{2p_1p_2} \right) + \gamma \cdot \frac{2^{1/3}}{2^{1/3} - 1} \frac{1}{8} \frac{n^3}{p} + \mathcal{O}(\alpha \log^2 p).$$

In contrast to LU factorization and our recursive TRSM algorithm, this is a much lower synchronization cost.

VI. ITERATIVE TRIANGULAR SOLVER

In this section, we present our main contribution, a 3D TRSM algorithm that uses inversion of diagonal blocks to achieve a lower synchronization cost. By precomputing the inversions, we replace the latency-dominated small TRSMs with more parallel matrix multiplications.

A. Block-Diagonal Triangular Inversion

In order to lower the synchronization cost of TRSM, we start by inverting a set of triangular blocks along the diagonal of the matrix, each with a distinct subset of processors. We achieve this by splitting the p processors of the grid into $\frac{n}{n_0}$ different subgrids and each of those subgrids gets assigned one block inversion. To have the proper layout for the inversion, a transition from the original, block cyclic layout on a subgrid to the grid as described in Section III has to happen. With this layout, all the inversions can be done in parallel. The dimensions of each of the subgrids is $r_1 \times r_1 \times r_2$, where $r_1^2 r_2 = p \frac{n_0}{n}$. To support our inversion, we must have $r_2 > r_1$ and $n_0 \geq \sqrt{r_1^2 r_2}$. The precise choices of

r_1 and r_2 are given in the algorithm and will be discussed in Section VII.

$$\tilde{L} = \text{Diagonal-Inverter}(L, \Pi, n, p_1, p_2, n_0)$$

Require:

The processor grid Π has dimensions $p_1 \times p_1 \times p_2$
 L is a lower triangular $n \times n$ matrix distributed cyclically on Π
such that processor $\Pi(x, y, 1)$ owns $L[x, y]$ a lower triangular
 $\frac{n}{p_1} \times \frac{n}{p_1}$ matrix such that $L[x, y](i, j) = L(ip_1 + x, jp_1 + y)$.

- 1: Define $q = \frac{pn_0}{n}$ $r = \frac{n}{n_0}$
 - 2: Define $r_1 = \left(\frac{pn_0}{4n}\right)^{1/3}$
 - 3: Define $r_2 = \left(\frac{16pn_0}{n}\right)^{1/3}$
 - 4: Define a $p_1 \times p_1 \times \sqrt{p_2} \times \sqrt{p_2}$ processor grid Π_{4D} , such that
 $\Pi_{4D}(x_1, x_2, y_1, y_2) = \Pi(x_1, x_2, y_1 + \sqrt{p_2}y_2)$.
and $\Pi_{4D}(x_1, x_2, y_1, y_2)$ owns blocks $L[x_1, x_2, y_1, y_2]$
 - 5: Define a block diagonal matrix $L_D[x_1, x_2, y_1, y_2]$
such that $L_D[x_1, x_2, y_1, y_2][b]$ denotes
the block $L[x_1, x_2, y_1, y_2](bn_0 : (b+1)n_0, bn_0 : (b+1)n_0)$
 - 6: **Scatter** ($L_D[\circ, \circ, y_1, y_2][\circ]$,
 $\Pi_{4D}(x_1, x_2, 1, 1), \Pi_{4D}(x_1, x_2, \circ, \circ)$)
 - 7: Define a $\sqrt{p} \times \sqrt{p}$ processor grid Π_{2D} , such that
 $\Pi_{2D}(x_1 + p_1y_1, x_2 + p_2y_2) = \Pi_{4D}(x_1, x_2, y_1, y_2)$.
 - 8: Define a $\sqrt{q} \times \sqrt{q} \times \sqrt{r} \times \sqrt{r}$ processor grid Π_{4D}^I , such that
 $\Pi_{4D}^I(u_1, u_2, v_1, v_2) = \Pi_{2D}(u_1 + \sqrt{q}v_1, u_2 + \sqrt{q}v_2)$ owns
blocks $L_D[u_1, u_2, v_1, v_2]$.
 - 9: **AllToAll** ($L_D[u_1, u_2, v_1, v_2][\circ]$, $\Pi_{4D}^I(u_1, u_2, \circ, \circ)$)
 - 10: **For** $i = 0 : \sqrt{\frac{n}{n_0}} - 1$ **do in parallel**
 - 11: **For** $j = 0 : \sqrt{\frac{n}{n_0}} - 1$ **do in parallel**
 - 12: Define $b = \left(i + \sqrt{\frac{n}{n_0}}j\right)$
 - 13: $\tilde{L}_D[\circ, \circ, i, j][b] = \text{RecTriInv}(L_D[\circ, \circ, i, j][b],$
 $\Pi_{4D}^I[\circ, \circ, i, j], q, r_1, r_2)$
 - 14: **end for**
 - 15: **end for**
 - 16: **AllToAll** ($\tilde{L}_D[u_1, u_2, v_1, v_2][\circ]$, $\Pi_{4D}^I(u_1, u_2, \circ, \circ)$)
 - 17: **Gather** ($\tilde{L}_D[\circ, \circ, y_1, y_2][\circ]$,
 $\Pi_{4D}(x_1, x_2, \circ, \circ), \Pi_{4D}(x_1, x_2, 1, 1)$)
 $\tilde{L}_D L_D = \mathbf{1} \quad \forall i$ where L and \tilde{L} are partitioned the same way
-

B. Triangular Solve using Partial Inversion

As initial condition for the algorithm we set L to be distributed on the top level of the three dimensional grid Π in a cyclic layout such that each processor $\Pi(x, y, 1)$ owns $L[y : p_1 : n, x : p_1 : n]$. Also, we set the right hand side to be distributed on one level of the grid with a blocked layout with a physical block size of $b \times \frac{k}{p_2}$ such that each processor $\Pi(x, 1, z)$ owns $B[x : p_1 : \frac{n}{b}, zk/p_2 : (z+1)k/p_2]$. Additionally, each processor has memory of the same size as its part of B allocated for an update-matrix denoted as \bar{B}_j , $j \in [1, p_1]$, where also each processor $\Pi(x, y, z)$ owns $\bar{B}_y[x : p_1 : \frac{n}{b}, zk/p_2 : (z+1)k/p_2]$. The algorithm itself consists of two parts: first “inversion”, we invert all the

base-cases on the diagonal in parallel as described in the algorithm above and, second “solve”, we do the updates and calculate the solution to TRSM.

$$X = \text{It-Inv-TRSM}(L, B, \Pi, n, k, p_1, p_2, r_1, r_2)$$

Require:

The processor grid Π has dimensions $p_1 \times p_1 \times p_2$
 L is a lower triangular $n \times n$ matrix is distributed on Π
such that $\Pi(x, y, 1)$ owns $L[x, y] \in \mathbb{R}^{n/p_1 \times n/p_1}$ such that
 $L[x, y](i, j) = L(ip_1 + x, jp_1 + y)$
 B is a dense $n \times k$ matrix is distributed such that $\Pi(x, 1, z)$
owns $B[x, z] \in \mathbb{R}^{n/p_1 \times k/p_2}$, such that $B[x, z](i, j) =$
 $L(ip_1 + x, zk/p_2 + j)$
Define the block $S_i = in_0 : (i+1)n_0$ and the block $T_i =$
 $in_0 : n$

- 1: $\tilde{L} = \text{Diagonal-Inverter}(L, \Pi, n, p_1, p_2, n_0)$
- 2: **Bcast** ($B[x, z](S_0(x), \circ), \Pi(x, 1, z), \Pi(x, \circ, z)$)
- 3: **for** $i = 0 : \frac{n}{n_0} - 1$ **do**
- 4: $\Pi(x, y, z) : X[y, z](S_i, \circ) =$
 $\tilde{L}[y, x](S_i, S_i) \cdot B[x, z](S_i, \circ)$
- 5: $X[y, z](S_i, \circ) = \text{Allreduce}(X[y, z](S_i, \circ), \Pi(\circ, y, z))$
- 6: **Bcast** ($\tilde{L}[x, y](T_{i+1}, S_i), \Pi(x, y, 1), \Pi(x, y, \circ)$)
- 7: $\Pi(x, y, z) : \bar{B}_y[x, z](T_{i+1}, \circ) =$
 $\tilde{L}[x, y](T_{i+1}, S_i) \cdot X[y, z](S_i, \circ)$
- 8: $\bar{B}_0[x, z](S_{i+1}, \circ) =$
 $\text{Allreduce}(\bar{B}_\circ[x, z](S_{i+1}, \circ), \Pi(x, \circ, z))$
- 9: $\Pi(x, y, z) : B[x, z](S_{i+1}, \circ) =$
 $B[x, z](S_{i+1}, \circ) - \bar{B}_0[x, z](S_{i+1}, \circ)$
- 10: **end for**

Ensure:

$B = LX$ where X is distributed the same way as B

VII. COST ANALYSIS OF THE ITERATIVE TRSM

In this section we will derive a performance model for the algorithm presented in Section VI. The total cost of the algorithm is put together from the cost of its three subroutines

$$T_{\text{It-Inv-TRSM}}(n, k, n_0, p_1, p_2) = T_{\text{Inv}}(n, p_1, p_2) +$$

$$T_{\text{Upd}}(n, k, n_0, p_1, p_2) + T_{\text{Solve}}(n, k, n_0, p_1, p_2),$$

where the cost denoted by inversion is the part of the algorithm that inverts the blocks (Algorithm Diagonal-Inverter). The solve part is in lines 4-5, and the update in lines 7-8.

A. Inversion Cost

We invert the $\frac{n}{n_0}$ submatrices of size $n_0 \times n_0$ along the diagonal with distinct processor grids. The size of the processor grids involved is $r_1 \times r_1 \times r_2$. The choices of r_1 and r_2 are made such that the bandwidth cost of the inversion is minimal. Additionally we have to account for the cost that arises from communicating the submatrices to the proper subgrids. This happens in lines 6, 9, 16, and 17

of the Algorithm Diagonal-Inverter. The respective costs are summed up the the following table:

Line 6	$\alpha \cdot \log(p_2) + \beta \cdot \frac{nn_0}{2p_1^2}$
Line 9	$\mathcal{O}\left(\alpha \cdot \log(p) + \beta \cdot \frac{nn_0 \log p}{2p}\right)$
Line 16	$\mathcal{O}\left(\alpha \cdot \log(p) + \beta \cdot \frac{nn_0 \log p}{2p}\right)$
Line 17	$\alpha \cdot \log(p_2) + \beta \cdot \frac{nn_0}{2p_1^2}$

These cost are never of leading order compared to the costs that arise from the triangular inversion. With the derivations done in Section V, we get the following costs for inversion:

Latency Cost: The total latency cost of inversion is

$$S_{\text{Inv}}(p) = \mathcal{O}(\alpha \log^2 p).$$

Bandwidth Cost: In order to minimize the bandwidth cost of triangular inversion, we choose a grid splitting to achieve closest to ideal ratios for the subgrids processor layout r_1 and r_2 . This ratio is achieved when $r_2 = 4r_1$. The choices for r_1 and r_2 are therefore

$$r_1 = \left(\frac{pn_0}{4n}\right)^{1/3} \quad \text{and} \quad r_2 = \left(\frac{16pn_0}{n}\right)^{1/3}.$$

With this grid slicing, we get $\frac{n}{n_0}$ different sub-grids of dimensions $r_1 \times r_1 \times r_2$. This setup leads to a cost for inverting $\frac{n}{n_0}$ submatrices of $W_{\text{Inv}}(n_0, r_1, r_2) =$

$$W_{\text{RecTriInv}}(n_0, r_1, r_2) = \frac{2^{1/3}}{2^{1/3} - 1} \left(\frac{n_0^2}{8r_1^2} + \frac{n_0^2}{2r_1 r_2} \right).$$

Flop Cost: The flop cost of the inversion part is

$$F_{\text{Inv}}(n_0, p_1, p_2) = \frac{1}{8} \frac{nn_0^2}{p_1^2 p_2}.$$

B. Solve Cost

The complete solve cost can be derived by

$$T_{\text{Solve}}(n, n_0, k, p, p_1, p_2) = \frac{n}{n_0} T_{\text{MM}}(n_0, k, p, p_1, p_2)$$

Latency Cost: The latency cost of the solve part is

$$S_{\text{Solve}}(n, n_0, p) = \mathcal{O}\left(\frac{n}{n_0} \log p\right).$$

Bandwidth Cost: The cost of the solve is one call to triangular matrix multiplication for each base case. The synchronization cost is again dominated by the $\frac{n}{n_0}$ cases. The cost for these sums has been presented in Section III. Including these, we obtain a total cost of

$$\begin{aligned} W_{\text{Solve}}(n, k, n_0, p, p_1, p_2) &= \frac{n}{n_0} \cdot W_{\text{MM}}(n_0, k, p, p_1, p_2) \\ &= \frac{n}{n_0} \cdot \left[\left(\frac{n_0^2}{p_1^2} \right) \mathbb{1}_{p_2} + 4 \left(\frac{n_0 k}{p_1 p_2} \right) \mathbb{1}_{p_1} \right]. \end{aligned}$$

Flop Cost: The flop cost of the solve part is

$$F_{\text{Solve}}(n, k, n_0, p_1, p_2) = \frac{n}{n_0} \left(\frac{n_0^2 k}{p_1^2 p_2} \right).$$

C. Update Cost

The complete solve cost can be derived by

$$T_{\text{Upd}}(n, k, n_0, p, p_1, p_2) = \sum_{i=1}^{n/n_0-1} T_{\text{MM}}(n - in_0, n_0, k, p, p_1, p_2).$$

Latency Cost: The update latency cost is

$$S_{\text{Upd}}(n, n_0, p) = \mathcal{O}\left(\frac{n - n_0}{n_0} \log p\right).$$

Bandwidth Cost: The cost of doing all the updates as described in the algorithm in Section VI (Lines 5-9) is the cost of both the allreductions and the broadcast,

$$\begin{aligned} W_{\text{Upd}}(n, k, n_0, p_1, p_2) &= \sum_{i=1}^{n/n_0-1} \left[W_{\text{broadcast}}\left(\frac{nn_0 - in_0}{p_1^2}, p_2\right) + \right. \\ &\quad \left. W_{\text{allreduction}}\left(\frac{n_0 k}{p_1 p_2}, p_1\right) + W_{\text{allreduction}}\left(\frac{n_0 k}{p_1 p_2}, p_1\right) \right]. \end{aligned}$$

This yields to a total cost of $W_{\text{Upd}}(n, k, n_0, p_1, p_2) =$

$$\frac{n - n_0}{n_0} \left[4 \frac{nn_0 - n}{p_1^2} \mathbb{1}_{p_2} + 4 \frac{n_0 k}{p_1 p_2} \mathbb{1}_{p_1} \right].$$

Flop Cost: The update flop cost is

$$F_{\text{Upd}}(n, k, n_0, p_1, p_2) = \frac{n - n_0}{n_0} \left(\frac{kn n_0}{p_1^2 p_2} \right).$$

D. Total Cost

The total cost of the Algorithm is the sum of its three parts and leaves a lot of tuning room as with a choice of $p_1 = 1$, $p_2 = 1$ or $n_0 = n$ one is able to eliminate certain terms.

Latency Cost: The total latency cost of the algorithm is a sum of the previous parts,

$$\begin{aligned} S_{\text{It-Inv-TRSM}}(p_1, p_2, r_1, b) &= S_{\text{Upd}}(n, n_0, p_1, p_2) + \\ &\quad S_{\text{Solve}}(n, n_0, p_1, p_2) + S_{\text{Inv}}(p_1, p_2, r_1, b) \\ &= \mathcal{O}\left(\alpha \left(\frac{n}{n_0} \log p + \log^2 p\right)\right). \end{aligned}$$

Bandwidth Cost: The total bandwidth cost for the TRSM algorithm is, by choosing $\nu = \frac{2^{1/3}}{2^{1/3}-1}$,

$$\begin{aligned} W_{\text{It-Inv-TRSM}}(n, k, n_0, p_1, p_2, u, v, b) &= \\ &\quad W_{\text{Upd}}(n, k, n_0, p_1, p_2) + W_{\text{Solve}}(n, k, n_0, p_1, p_2) \\ &\quad + W_{\text{Inv}}(n, b, p_1, p_2, u, v) \\ &= \frac{n}{n_0} \cdot \left[\left(\frac{n_0^2}{p_1^2} \right) \mathbb{1}_{p_2} + 4 \left(\frac{n_0 k}{p_1 p_2} \right) \mathbb{1}_{p_1} \right] \\ &\quad + \frac{n - n_0}{n_0} \left[4 \frac{nn_0 - n}{p_1^2} \mathbb{1}_{p_2} + 4 \frac{n_0 k}{p_1 p_2} \mathbb{1}_{p_1} \right] + \nu \left(\frac{n_0^2}{8r_1^2} + \frac{n_0^2}{2r_1 r_2} \right). \end{aligned}$$

Flop Cost: Lastly, the combined total flop cost is

$$\begin{aligned} F_{\text{It-Inv-TRSM}}(n, k, n_0, p_1, p_2) &= F_{\text{Upd}}(n, n_0, p_1, p_2) \\ &\quad + F_{\text{Solve}}(n, n_0, p_1, p_2) + F_{\text{Inv}}(n_0, u, v, p_1, p_2) = \frac{n^2 k}{p_1^2 p_2} + \frac{n_0^2 n}{p_1^2 p_2} \end{aligned}$$

VIII. PARAMETER TUNING

In this section, we give asymptotically optimal tuning parameters for different relative matrix sizes to optimize performance. We only focus on asymptotic parameters as there is a trade off between the constant factors on the bandwidth and latency costs. The exact choice is therefore machine dependent and should be determined experimentally. The initial grid layout is dependent on the relative matrix sizes of L and B since the update part of the algorithm is one of the dominating terms in any case where there is an update to be made and determines the case where it is infeasible.

In the case of one large dimension ($n < \frac{4k}{p}$), the initial grid layout is $p_1 = 1$ and $p_2 = p$. The optimal parameters are given in the following table:

p_1	1	p_2	p	n_0	n
r_1	$\mathcal{O}\left((p)^{1/3}\right)$	r_2	$\mathcal{O}\left((p)^{1/3}\right)$		

Using this set of parameters will yield to a total cost of

$$T_{\text{IT1D}}(n, k, p) = \mathcal{O}\left(\alpha \cdot (\log^2 p + \log p) + \beta \cdot n^2 + \gamma \cdot \frac{n^2 k}{p}\right).$$

Comparing these costs to the costs of T_{RT1D} obtained in Section IV-A, we can see that we are within the asymptotic bounds of the original algorithm in bandwidth and flop cost, but pay an extra factor of $\log p$ latency, since the inversion, if performed on a 3D grid, requires $\log^2 p$ steps. But since the inversion is the least significant part of the routine in 1 large dimension, no gain was to be expected in this case.

In the case of two large dimensions ($n > 4k\sqrt{p}$), the initial grid layout is $p_1 = \sqrt{p}$ and $p_2 = 1$. The optimal parameters are given in the following table:

p_1	\sqrt{p}	p_2	1	n_0	$\mathcal{O}\left(\left(nk^3 p^{1/2}\right)^{1/4}\right)$
r_1	$\mathcal{O}\left(\left(\frac{k}{n}\right)^{1/4} p^{3/8}\right)$	r_2	$\mathcal{O}\left(\left(\frac{k}{n}\right)^{1/4} p^{3/8}\right)$		

Using this set of parameters will yield to a total cost of

$$T_{\text{IT2D}}(n, k, p) = \mathcal{O}\left(\alpha \left(\log^2 p + \left(\frac{n}{k}\right)^{3/4} \frac{1}{p^{1/8}} \log p\right) + \beta \left(\frac{nk}{\sqrt{p}}\right) + \gamma \left(\frac{n^2 k}{\sqrt{p}}\right)\right).$$

Comparing these costs to the cost of T_{RT2D} obtained in Section IV-A, we can see that we are asymptotically more efficient in terms of latency by a factor of at least $\frac{p^{1/4}}{\log p}$ as well as in bandwidth by a factor of $\log p$ while being able to keep the flop costs asymptotically constant. This is a significant gain and especially important as the occurrence of fewer right hand sides $k < n$ is high.

In the case of three large dimensions $\frac{4k}{p} \leq n \leq 4k\sqrt{p}$, the initial grid layout is $p_1 = p^{1/3} \left(\frac{n}{4k}\right)^{1/3}$ and $p_2 = p^{1/3} \left(\frac{4k}{n}\right)^{2/3}$. The optimal parameters are given in the following table:

p_1	$\left(\frac{pn}{4k}\right)^{1/3}$	p_2	$\left(\frac{\sqrt{p}4k}{n}\right)^{2/3}$	n_0	$\mathcal{O}\left(\min\left(\sqrt{nk}, n\right)\right)$
r_1	$\mathcal{O}\left(\left(\min\left[\frac{p\sqrt{nk}}{n}, p\right]\right)^{1/3}\right)$	r_2	$\mathcal{O}\left(\left(\min\left[\frac{p\sqrt{nk}}{n}, p\right]\right)^{1/3}\right)$		

Using this set of parameters will yield to a total cost of

$$T_{\text{IT3D}}(n, k, p) = \mathcal{O}\left(\alpha \left(\log^2 p + \max\left(\sqrt{\frac{n}{k}}, 1\right) \log p\right) + \beta \left(\left(\frac{n^2 k}{p}\right)^{2/3}\right) + \gamma \left(\frac{n^2 k}{\sqrt{p}}\right)\right).$$

Comparing these costs to the cost of T_{RT3D} obtained in Section IV-A, we can see that we are asymptotically more efficient in terms of latency by a factor of $\left(\frac{n}{k}\right)^{1/6} p^{2/3}$ while being able to keep bandwidth and flop costs asymptotically constant.

IX. CONCLUSION

	S	W	F
	1 Large Dimension ($n < \frac{4k}{p}$)		
standard	$\log p$	n^2	$\frac{n^2 k}{p}$
new method	$\log^2 p$	n^2	$\frac{n^2 k}{p}$
	2 Large Dimensions ($n > 4k\sqrt{p}$)		
standard	\sqrt{p}	$\log p \frac{nk}{\sqrt{p}}$	$\frac{n^2 k}{p}$
new method	$\log^2 p + \left(\frac{n}{k}\right)^{3/4} \frac{1}{p^{1/8}} \log p$	$\frac{nk}{\sqrt{p}}$	$\frac{n^2 k}{p}$
	3 Large Dimensions $\frac{4k}{p} \leq n \leq 4k\sqrt{p}$		
standard	$\left(\frac{np}{k}\right)^{2/3} \log p$	$\left(\frac{n^2 k}{p}\right)^{2/3}$	$\frac{n^2 k}{p}$
new method	$\log^2 p + \sqrt{\frac{n}{k}} \log p$	$\left(\frac{n^2 k}{p}\right)^{2/3}$	$\frac{2n^2 k}{p}$

We presented a new method for solving triangular systems for multiple right hand sides. Our algorithm was able to achieve better theoretical scalability than known alternatives by up to a factor of $\left(\frac{n}{k}\right)^{1/6} p^{2/3}$. For certain matrix dimensions, a decrease of bandwidth cost by a factor of $\log_2 p$ was obtained. All the asymptotic costs obtained with the new algorithm are shown in the table above. This was achieved by a selective use of triangular matrix inversion. A recursive approach for inverting triangular matrices was used where a high amount of parallelism and a low latency was achieved. By only inverting triangular blocks along the diagonal of the initial matrix, we generalized the usual way of TRSM computation and the full matrix inversion approach. The flexibility of fine tuning the algorithm based on the relative input sizes as well as the number of processors available led to an efficient algorithm. The cost analysis of this new method allowed us to give recommendations for asymptotically optimal tuning parameters for a wide variety of possible inputs. With the detailed pseudo code provided, we were able to discuss communication patterns in depth and make an implementation as easy as possible.

REFERENCES

- [1] F. G. Gustavson, "Recursion leads to automatic variable blocking for dense linear-algebra algorithms," *IBM Journal of Research and Development*, vol. 41, no. 6, pp. 737–755, 1997.
- [2] E. Solomonik and J. Demmel, "Communication-optimal parallel 2.5 D matrix multiplication and LU factorization algorithms," in *Euro-Par 2011 Parallel Processing*. Springer, 2011, pp. 90–109.
- [3] E. Elmroth, F. Gustavson, I. Jonsson, and B. Kågström, "Recursive blocked algorithms and hybrid data structures for dense matrix library software," *SIAM review*, vol. 46, no. 1, pp. 3–45, 2004.
- [4] J. J. Du Croz and N. J. Higham, "Stability of methods for matrix inversion," *IMA Journal of Numerical Analysis*, vol. 12, no. 1, pp. 1–19, 1992.
- [5] E. Solomonik, E. Carson, N. Knight, and J. Demmel, "Trade-offs between synchronization, communication, and computation in parallel linear algebra computations," in *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures*, ser. SPAA '14. New York, NY, USA: ACM, 2014, pp. 307–318.
- [6] E. Chan, M. Heimlich, A. Purkayastha, and R. Van De Geijn, "Collective communication: theory, practice, and experience," *Concurrency and Computation: Practice and Experience*, vol. 19, no. 13, pp. 1749–1783, 2007.
- [7] R. Thakur, R. Rabenseifner, and W. Gropp, "Optimization of collective communication operations in mpich," *International Journal of High Performance Computing Applications*, vol. 19, no. 1, pp. 49–66, 2005.
- [8] J. Bruck, C.-T. Ho, S. Kipnis, E. Upfal, and D. Weatherby, "Efficient algorithms for all-to-all communications in multiport message-passing systems," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 8, no. 11, pp. 1143–1156, 1997.
- [9] J. L. Träff and A. Ripke, "Optimal broadcast for fully connected processor-node networks," *Journal of Parallel and Distributed Computing*, vol. 68, no. 7, pp. 887–901, 2008.
- [10] E. Dekel, D. Nassimi, and S. Sahni, "Parallel matrix and graph algorithms," *SIAM Journal on Computing*, vol. 10, no. 4, pp. 657–675, 1981.
- [11] R. C. Agarwal, S. M. Balle, F. G. Gustavson, M. Joshi, and P. Palkar, "A three-dimensional approach to parallel matrix multiplication," *IBM J. Res. Dev.*, vol. 39, pp. 575–582, September 1995.
- [12] A. Aggarwal, A. K. Chandra, and M. Snir, "Communication complexity of PRAMs," *Theoretical Computer Science*, vol. 71, no. 1, pp. 3 – 28, 1990.
- [13] J. Berntsen, "Communication efficient matrix multiplication on hypercubes," *Parallel Computing*, vol. 12, no. 3, pp. 335–342, 1989.
- [14] W. F. McColl and A. Tiskin, "Memory-efficient matrix multiplication in the BSP model," *Algorithmica*, vol. 24, pp. 287–297, 1999.
- [15] S. L. Johnsson, "Minimizing the communication time for matrix multiplication on multiprocessors," *Parallel Comput.*, vol. 19, pp. 1235–1257, November 1993.
- [16] J. Demmel, D. Eliahu, A. Fox, S. Kamil, B. Lipshitz, O. Schwartz, and O. Spillinger, "Communication-optimal parallel recursive rectangular matrix multiplication," in *Parallel Distributed Processing (IPDPS), 2013 IEEE 27th International Symposium on*, May 2013, pp. 261–272.
- [17] A. Aggarwal, A. K. Chandra, and M. Snir, "Communication complexity of prams," *Theoretical Computer Science*, vol. 71, no. 1, pp. 3–28, 1990.
- [18] A. Tiskin, "Bulk-synchronous parallel gaussian elimination," *Journal of Mathematical Sciences*, vol. 108, no. 6, pp. 977–991, 2002.
- [19] M. T. Heath and C. H. Romine, "Parallel solution of triangular systems on distributed-memory multiprocessors," *SIAM Journal on Scientific and Statistical Computing*, vol. 9, no. 3, pp. 558–588, 1988.
- [20] D. Irony and S. Toledo, "Trading replication for communication in parallel distributed-memory dense solvers," *Parallel Processing Letters*, vol. 12, no. 01, pp. 79–94, 2002.
- [21] B. Lipshitz, "Communication-avoiding parallel recursive algorithms for matrix multiplication," Master's thesis, University of California, Berkeley, 2013.
- [22] A. Borodin and I. Munro, *The computational complexity of algebraic and numeric problems*. Elsevier Publishing Company, 1975, vol. 1.
- [23] S. M. Balle, P. C. Hansen, and N. Higham, "A Strassen-type matrix inversion algorithm," *Advances in Parallel Algorithms*, pp. 22–30, 1994.